

Migration of User Home Directories From AFS To NAS (NFS/CIFS)

Andy Romero

High Level Summary

- ▶ Home directories and their contents are being migrated
 - ▶ FROM: The AFS user home directory areas
 - ▶ /afs/fnal.gov/files/home/room1
 - ▶ /afs/fnal.gov/files/home/room2
 - ▶ /afs/fnal.gov/files/home/room3
 - ▶ TO: One of the NAS based home directory areas
 - ▶ The UNIX home directory area
 - ▶ The Archive home directory area
 - ▶ In addition to the UNIX and Archive home directory areas, there is also a Windows/Mac home directory area. This document will discuss all three home directory areas; however, the Windows/Mac home directory area will not be a destination when migrating files from AFS.
- ▶ Sample migration source and destination paths
 - ▶ User hugo is an active employee
 - ▶ Source: /afs/fnal.gov/files/home/room3/hugo
 - ▶ Destination (nfs path): homesrv01.fnal.gov:/home/h/hugo
 - ▶ User waldo is retired (and is no longer working at the lab)
 - ▶ Source: /afs/fnal.gov/files/home/room2/waldo
 - ▶ Destination (nfs path): homesrv01.fnal.gov:/home-archive/w/waldo/afs-archive

NAS Based Home Directory Areas

- ▶ UNIX home directories
 - ▶ Intended Users: Scientists and other serious UNIX users.
 - ▶ Standard Posix style security model
 - ▶ AFS home directories for active users will be migrated here
 - ▶ There are about 2000 active AFS home directories
- ▶ Windows / Mac home directories
 - ▶ Intended Users: People who use Windows and Mac systems often; but, who never or hardly ever use Linux systems.
 - ▶ Microsoft Windows style ACL security model (this is the expected security model for CIFS/SMB clients)
 - ▶ No AFS home directories will be migrated to the Windows / Mac home directory area; however, it will possible for non-Unix users to easily access their UNIX home directories (and their migrated legacy AFS files) from both MacOS and Windows systems.
- ▶ Archive home directories
 - ▶ AFS home directories for inactive users will be migrated here
 - ▶ There are over 5000 inactive AFS home directories

UNIX Home Directories

- ▶ **Unix Home Directory Access Path**
 - ▶ NFS Volume Path.....: homesrv01.fnal.gov:/home
 - ▶ On gpcf/gpvm and fnalu nodes the NFS volume homesrv01.fnal.gov:/home will be mounted as
 - ▶ /nashome
 - ▶ The path of an individual user's home directory on these nodes will be of the form: /nashome/roomID/username where roomID is one of the following letters {a,b,c, ..., z}. The roomID for a user's home directory will be the first letter of the user's username
 - ▶ Sample home directory paths for users
 - ▶ /nashome/s/smith
 - ▶ /nashome/j/jones
 - ▶ The NAS administrators recommend that other system admins also mount the NFS volume homesrv01.fnal.gov:/home using the local path /nashome
(detailed mounting instructions can be found in the File Server Guide For Linux Users and Administrators)
- ▶ **Trusted Hosts can mount homesrv01.fnal.gov:/home using NFSv3**
 - ▶ What is a trusted host ?
 - ▶ Answer: It is a host that has all of the following characteristics
 - ▶ Managed by the Scientific Server Team
 - ▶ End users must present Kerberos credentials to logon
 - ▶ NIS DB and/or password / group files are under the exclusive control of the official Scientific Server team sys-admins
 - ▶ Located in a secure computer room
- ▶ **All other Linux hosts must mount homesrv01.fnal.gov:/home using NFSv4-Krb**
 - ▶ This includes
 - ▶ Linux Servers not managed by Scientific Server Team
 - ▶ Linux Workstations

UNIX Home Directory Security

- ▶ POSIX security configuration on the top level of each individual user's home directory.
 - ▶ The UNIX Owner will be set to the end-user's UID
 - ▶ The UNIX Group will be set to GID=3000 (nas-admin)
This is a safe group; it will have no members defined in the NAS (server-side) group store or in the (client-side) Trusted Host NIS DBs
 - ▶ The setgid bit will be set; so that, the safe group (GID=3000) is the default group of all new items created. This prevents the accidental granting of unwanted access authorization to new items. Users can use chgrp to explicitly control the group of new files and subdirectories they create. Of course for a user to "chgrp" a file/sub-directory to a group, the user must be a member of that group.
 - ▶ The Permission bits will be set to 711
These permissions allow for easy sharing of data with co-workers who know the full path of the item being shared; however, these permissions prevent "rm -rf" attacks; which, if launched from /nashome, would traverse down into every user's home directory and annihilate all misconfigured sub-directories and files. 711 permissions also prevent unwanted file snooping.
 - ▶ The storage managers will enforce the 711 permissions on the root (top level) of each individual user's home directory. Users will be able to use standard UNIX tools (umask, chmod, chgrp ..) to fully control the permissions of the items they create under their home directory.

Use of Shared Accounts (Shared UID)

- ▶ What is a Shared Account / Shared UID
 - ▶ If two or more different Kerberos identities (principals) can run processes on a host using a single UID identity, then that UID is shared.
- ▶ Examples of shared Accounts / UIDs
 - ▶ Experiment controlled general purpose shared accounts: example: lbne-daq
 - ▶ The Fermilab cron principals designed for use with Fermilab's Customized AFS implementation. (A user's main principal and a user's Fermi cron principal are 2 separate Kerberos identities sharing a common UID identity)
- ▶ The Linux NFSv4 client currently does not support multiple Krb identities being associated with a single local UID identity. The Kernel portion of the Linux NFS client associates only one Kerberos identity with all processes run in the context of the shared UID.

Use of Shared Accounts (Shared UID)

- ▶ Use of Shared Accounts, on hosts which mount home directory volumes, is only supported on Trusted Hosts.
 - ▶ A word on cron principals: Fermi personal cron principals can be used on Trusted Hosts; however, they are superfluous on Trusted Hosts.
- ▶ Untrusted hosts (workstations ...etc), on which shared accounts are used in a non-trivial manner, should not mount the home directory volumes.
 - ▶ “indirect” access methods (scp) can be used to access personal home directory files from un-trusted hosts which make non-trivial use of shared accounts.
- ▶ Long running jobs, run on un-trusted hosts, (that need to access NFSv4 based files) should:
 - ▶ Be run using a non-personal, task dedicated “robot” Kerberos identity/principal which should have its own official UID.
 - ▶ Have the robot identity added to the server-side group that controls access to the relevant NFSv4 hosted directory.

UNIX Home Directories

..access from Windows and Mac clients

- ▶ Summary

The UNIX home directories are configured to operate according to established UNIX/Linux norms. However; it is a fact that many Linux users occasionally use Windows and Mac systems and may need to exchange files between these systems and their Linux systems; therefore it will be possible for users to access their UNIX home directories from Windows and Mac systems via the CIFS/SMB protocol.
- ▶ CIFS / SMB Access paths
 - ▶ CIFS Path (win): <\\homesrv01\home-unix>
 - ▶ CIFS Path (Mac).....: smb://homesrv01.fnal.gov/home-unix
 - ▶ Authentication Domain for CIFS/SMB access: FERMI

Archived Home Directories

- ▶ NFS Volume Path: homesrv01.fnal.gov:/home-archive
 - ▶ No NFSv3 access NFSv4-krb only
- ▶ CIFS Path (win): \\homesrv01\home-archive
- ▶ CIFS Path (Mac): smb://homesrv01.fnal.gov/home-archive
- ▶ The home-archive volume is a read-only volume
- ▶ Who can access the home-directories in the home-archive volume
 - ▶ The NAS Admins
 - ▶ End users who return to the lab and re-gain their Kerberos credentials
 - ▶ Authorized supervisors can request that an in-active user's files be transferred. Requests will be reviewed and serviced on a case-by-case basis according to established "transfer of personal data" rules.

Windows Home Directories

- ▶ Windows Home Directory Volume Access Path
 - ▶ CIFS Path (win): \\homesrv01\users
 - ▶ CIFS Path (Mac).....: smb://homesrv01.fnal.gov/users
 - ▶ The path of an individual user's home directory will be of the form:
 - ▶ Windows: \\homesrv01\users\roomID\username
 - ▶ Mac: smb://homesrv01.fnal.gov/users/roomID/username
 - ▶ roomID is one of the following letters {a,b,c, ..., z}. The roomID for a user's home directory will be the first letter of the user's username
 - ▶ Sample home directory path for a user named smith
 - ▶ Win: \\homesrv01\users\s\smith
 - ▶ Mac: smb://homesrv01.fnal.gov/users/s/smith
- ▶ Authentication Domain: FERMI (fermi.win.fnal.gov)

Windows Home Directory Security

- ▶ Windows Home Directory Security is ACL based and unlike UNIX/Posix, the security settings for newly created items are not client umask driven; the security on new items is inherited from the parent container according to rules specified in the ACL.
- ▶ Individual user's home-directory-root (\\homesrv01\users\s\smith)
 - ▶ Owner (smith in this case): Full Access
 - ▶ Everyone: Traverse Only
(Traverse-only allows co-workers to access the special PUBLIC_FILE sub-directory; traverse-only does not grant any read, list, browse or write privilege)
- ▶ Files and subdirectories end-users create under their home-directory-root
 - ▶ Owner: Full Access
 - ▶ Others have no access (The everyone traverse-only permission is NOT inherited by new files and subdirectories created under a user's home-directory-root)
- ▶ Individual user's preconfigured PUBLIC_FILE subdirectory (\\homesrv01\users\s\smith\PUBLIC_FILE)
 - ▶ Owner: Full Access
 - ▶ Everyone: Traverse Only
- ▶ Files and subdirectories end-users create under their PUBLIC_FILE subdirectory
 - ▶ End User: Full Access
 - ▶ Everyone: Read-Only Access to files and sub-directories.
Since co-workers cannot list the contents of the PUBLIC_FILE subdirectory, the owner will need to give them the full path of the item being shared. This scheme provides users with a simple low overhead way to have multiple independent sharing relationships.

Windows Home Directories ...access from Linux

- ▶ **Summary**
The Windows home directories are configured to operate according to established Windows file server norms. However; it is a fact that some Windows / Mac users may occasionally use Linux systems and may need to exchange files between these systems and their Windows / Mac systems; therefore it will be possible for users to access their Windows home directories from Linux systems via the NFSv4-Krb protocol.
- ▶ **NFS Path.....: homesrv01.fnal.gov:/users**
 - ▶ No hosts will be granted NFSv3 access.
 - ▶ NFSv4 (with Kerberos) must be used
- ▶ **On fnalu nodes the NFS volume homesrv01.fnal.gov:/users is mounted as**
 - ▶ /naswinusers
 - ▶ The path of an individual user's Windows home directory on these nodes will be of the form: /naswinusers/roomID/username where roomID is one of the following letters {a,b,c, ..., z}. The roomID for a user's Windows home directory will be the first letter of the user's username
 - ▶ Sample home directory path for a user
 - ▶ /naswinusers/s/smith
- ▶ The NAS administrators recommend that when other system admins mount the NFS volume homesrv01.fnal.gov:/users , they also use the local path /naswinusers

Accessing Home Directories from offsite

- ▶ Initially there will be 2 options for exchanging files between offsite hosts and the NAS file service.
 - ▶ Connect to the lab via VPN
 - ▶ SCP through an onsite gpcf or fnalu host

Snapshots and Backups

- ▶ File system snapshots of the home directory areas
 - ▶ Rule based snapshots will be taken at 8AM, 10AM, 12PM, 2PM and 4PM
 - ▶ Snapshots will have a lifetime of 7 days
 - ▶ A backup driven snapshot will be taken later in the evening.
 - ▶ End users have self-service read-only access to snapshots via the .snapshot sub-directory located in the NFS volume root. For Win /Mac (CIFS / SMB) clients this is the ~snapshot directory located in the CIFS/SMB volume root.
- ▶ Nightly tape backups are performed using Tibs

Migration of AFS home directories

- ▶ AFS home directories for active users
 - ▶ The NAS / AFS admins will migrate files to the NAS based UNIX home directory area.
 - ▶ Permissions of migrated items will be secure by default; people other than the owner get no access to migrated files. The owner is granted rwx on migrated sub-directories and rw on migrated files unless the file had the execute mode bit set in AFS (in this case, to avoid breaking end user scripts we will preserve the execute bit state).
 - ▶ The default quota for the home directories in the NAS UNIX home directory store will be 2GB; however, additional quota will be granted as needed to facilitate the migration.
- ▶ AFS home directories for inactive users
 - ▶ The NAS / AFS admins will migrate files to the Archive home directory area.
 - ▶ If user waldo's files are to be archived, they will be transferred to `homesrv01.fnal.gov:/home-archive/w/waldo/afs-archive`.
 - ▶ It is possible that a future project will involve archiving the files, of former staff members, from other legacy servers. These would go into dedicated sub-directories under the user's home-archive directory ... example: `homesrv01.fnal.gov:/home-archive/w/waldo/cdserver-archive` .

Migration of AFS home directories

- ▶ The Migration process will be carried out in batches.
 - ▶ Early in the process (the “early adopter” phase), the batches (of active users) will be small (5-10users)
 - ▶ Batch size (for active users) will be progressively increased to as high as 500 after the change has been given “Go Live” approval.
 - ▶ Batch size, for the directories being archived will be much larger.

Migration of AFS home directories

- ▶ Workflow for a migration batch
 - ▶ Get batch user list from experiment liaisons (or from other sources: supervisors, Business Analysts etc.)
 - ▶ Communicate with the users assigned to the batch
 - ▶ Specify the cut-over date
 - ▶ Let them know AFS is your home directory until after cut-over
 - ▶ Other migration details (such as be logged out 30min before cut-overetc)
 - ▶ Keep source and destination in-sync to minimize cut-over downtime
 - ▶ One-day before cut-over, send out a reminder
 - ▶ Perform cut-over at specified time
 - ▶ Final file sync
 - ▶ Confirm sync was error free
 - ▶ Un-mount user's AFS home volume
 - ▶ Have Linux Admins make path changes to relevant NIS DBs or to password / group files
 - ▶ Users can logon and resume work.

Migration Responsibilities

- ▶ NAS / Storage Admins
 - ▶ Coordinate the process
 - ▶ Transfer / Sync the file payload
 - ▶ The NAS / Storage Admins will not be modifying the content of any user files.
- ▶ Linux Administrators
 - ▶ Manage the path changes to NIS / password / group stores
 - ▶ Update new-user on-boarding processes for their clusters
- ▶ The Users
 - ▶ Users who used hard-coded afs home directory path references in currently relevant scripts, links etc. are responsible for changing these to relative references, to ~ based references or to direct references to the new NAS based home directory.
The hard-coded afs path to a user's home directory will not be valid after cutover.

Post Migration Operations

- ▶ New Fermilab staff members will get a 2GB UNIX home directory and a 2GB Windows home directory.
 - ▶ Initially the workflow will involve a ticket to the NAS admins, who will create the new home directories.
 - ▶ A future goal is to have the new identity management system pass a home-directory creation request trigger file / token to the NAS management system which will automatically evaluate and process the request.
- ▶ Requests for quota increases will be via standard NAS/BlueArc storage request tickets; which can be submitted through the Fermilab Service Desk request system.

Additional Documentation

- ▶ File Server Guide For Linux Users and Administrators
- ▶ File Server Guide For Windows Users
- ▶ File Server Guide For Macintosh Users
- ▶ File Server Guide For IOS Users (coming soon)